

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 122 652 A1

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:

08.08.2001 Bulletin 2001/32

(51) Int Cl. 7: G06F 17/30

(21) Application number: 01102032.8

(22) Date of filing: 30.01.2001

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE TR

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 03.02.2000 US 497610

(71) Applicant: MITSUBISHI DENKI KABUSHIKI

KAISHA

Tokyo 100-8310 (JP)

- Young, Michael J.  
Boxborough, Massachusetts 01719 (US)
- DiCelle, Joseph J.  
Boylston, Massachusetts 01505 (US)
- Wong, David W. H.  
Boxborough, Massachusetts 01719 (US)
- Esenthaler, Alan W.  
Ashland, Massachusetts 01721 (US)

(74) Representative: Pfenning, Mehlig & Partner GbR  
Mozartstrasse 17  
80336 München (DE)

(72) Inventors:

- Walsh, Thomas C  
Cambridge, Massachusetts 02139 (US)

## (54) Data Integration system

(57) An enterprise integration system is coupled to a number of legacy data sources. The data sources each use different data formats and different access methods. The integration system includes a back-end interface configured to convert input data source information to input XML documents and to convert output XML document to output data source information. A front-end interface converts the output XML documents

to output HTML forms and the input HTML forms to the XML documents. A middle tier includes a rules engine and a rules database. Design tools are used to define the conversion and the XML documents. A network couples the back-end interface, the front-end interface, the middle tier, the design tools, and the data sources. Mobile agents are configured to communicate the XML documents over the network and to process the XML documents according to the rules.




FIG. 1b

Description	
<b>Field of the Invention</b>	
<p>[0001] This invention relates generally to computerized applications, databases, and interface, and more particularly to integrating applications, databases, and interfaces having different formats, contexts, and designs.</p>	discrete component manufacturing (such as in the automobile industry), and process manufacturing (such as chemical manufacturing through the use of real-time process control systems). Directives communicated from the computer systems to the manufacturing operations are commonly known as work orders. Work orders can include production orders, shipping orders, receiving orders, and the like.
<b>Background of the Invention</b>	
<p>[0002] Computer and computer-related technology have enabled the use of computers in numerous enterprise functions. Almost every facet of a modern enterprise is supported by computer systems in some manner. Computerization is a necessity to allow an enterprise to remain functional and competitive in a constantly changing environment.</p> <p>[0003] Computer systems are used to automate processes, to manage large quantities of information, and to provide fast and flexible communications. Many enterprises, from sole proprietorships, small stores, professional offices and partnerships, to large corporations have computerized their functions to some extent. Computers are pervasive, not only in business environment, but also in non-profit organizations, governments, and educational institutions.</p> <p>[0004] Computerized enterprise functions can include billing, order-taking, scheduling, inventory control, record keeping, and the like. Such computerization can be accomplished by using computer systems that run software packages. There are many application software packages available to handle a wide range of enterprise functions, including those discussed above.</p> <p>[0005] One such package is the SAP R/2™ System available from SAP America, Inc., 625 North Governor Printz Blvd., Essington, Pa. 19029. The SAP R/2 System is a software package designed to run on IBM or compatible mainframes in a CICS (Customer Interface Control System) or IMS (Information Management System) environment. For example, SAP may use CICS to interface with user terminals, printers, databases, or external communication facilities such as IBM's Virtual Telecommunications Access Method (VTAM).</p> <p>[0006] SAP is a modularized, table driven application software package that executes transactions to perform specified enterprise functions. These functions may include order processing, inventory control, and invoice validation; financial accounting, planning, and related managerial control; production planning and control; and project accounting, planning, and control. The modules that perform these functions are all fully integrated with one another.</p> <p>[0007] Another enterprise area that has been computerized is manufacturing. Numerous manufacturing functions are now controlled by computer systems. Such functions can include real-time process control of</p>	5 10 15 20 25 30 35 40 45 50 55

[0013] More than likely, the legacy system 10 will exhibit one or more of the following problems. All sub-systems cannot communicate with every other sub-system because each sub-system has its own application programming interfaces (APIs). Real-time data interchange among all of the sub-systems may be impossible or extremely difficult because each sub-system stores and views data in a different way and uses different communication protocols. Modified enterprise functions or adding automation for new functions is expensive. Each sub-system is developed with its own peculiar programming language. Users cannot always access all the data all of the time, particularly when the user is mobile. It is difficult to provide top level management with an abstraction of all system information.

[0014] What is needed is a system that can integrate various computer systems in an enterprise. The system needs to be able to convey transactional data between any number of databases regardless of their format, context, and access methodology. User interfaces to the databases need to be uniform. In addition, as enterprise functions change, new procedures and transactions must be accommodated in a minimal amount of time without having to redesign and reimplement any of the functional systems. The ideal enterprise integration system should be capable of adapting to any number of computerized functions in a modern complex enterprise.

#### **Summary of the Invention**

[0015] The present invention is directed to a system and method for integrating computer systems found in many types of enterprises.

[0016] An enterprise integration system is coupled to a number of legacy data sources. The data sources each use different data formats and different access methods. The integration system includes a back-end interface configured for converting input data source information to input XML documents and for converting output XML documents to output data source information.

[0017] A front-end interface converts the output XML documents to output HTML forms and the input HTML forms to the XML documents. A middle tier includes a rules engine and a rules database. Design tools are used to define the conversion and the XML documents.

[0018] A network couples the back-end interface, the front-end interface, the middle tier, the design tools, and the data sources. Mobile agents are configured to communicate the XML documents over the network and to process the XML documents according to the rules.

#### **Brief Description of the Drawings**

[0019]

Figure 1a is a block diagram of a legacy enterprise

system;

Figure 1b is a block diagram of an integrated enterprise system according to the invention;

Figure 2 is a block diagram of design tools used by the system of Figure 1b;

Figure 3 is a block diagram of XML data accesses according to the invention;

Figure 4 is a block diagram of a back-end interface of the system of Figure 1b;

Figure 5 is a diagrammatic of a public interface of the back-end interface of Figure 4;

Figure 6 is a block diagram of pooled connections;

Figure 7 is a flow diagram of a get request;

Figure 8 is a flow diagram of an update request; and

Figure 9 is a block diagram of an object of service bridge objects.

#### **Detailed Description of the Preferred Embodiments**

##### **Introduction**

[0020] Our invention provides a robust and scalable environment for integrating legacy enterprise computer systems. The invention integrates databases, transactions, and user interfaces having different formats, contexts, and designs, such as the sub-systems shown in Figure 1a. We also provide for automated rules based processing.

[0021] At the core of our integration system, we utilize XML as a universal data encoding and interchange format. XML (Extensible Markup Language) is a flexible way for us to create common information formats and share both the format and the data on the Internet, the World Wide Web (WWW), intranets, and private local area network. XML, developed by the World Wide Web

Consortium (W3C), is "extensible" because, unlike HyperText Markup Language (HTML), the markup symbols of XML are unlimited and self-defining. XML is actually a simpler and easier-to-use subset of the Standard Generalized Markup Language (SGML), the standard for how to create a document structure. XML enables us to create customized "tags" that provide functionality not available with HTML. For example, XML supports links that point to multiple documents, as opposed to HTML links, which can reference just one destination each.

[0022] These basic interfaces allow our integration system to view, modify and interact with linked legacy applications or legacy data sources.

## System Architecture

[0022] As shown in Figure 1b, our enterprise integration system 100 includes the following main components: a back-end interface 110, a front-end interface 120, a middle tier 130, and design tools 140. The components are connected by a network and mobile agents 101 carrying XML documents 102. The mobile agents 101 are described in greater detail in U.S. Patent Application Sn. 08/966, 716, filed by Walsh on November 7, 1997, incorporated herein in its entirety by reference. As a feature, the agents can travel according to itineraries, and agents can "meet" with each other at meeting points to interchange information.

[0023] With our back-end interface 110, we enable read/write/modify access to existing (legacy) applications and data sources 111. The back-end interface maps (or translates) data from legacy formats into the XML format used by our enterprise integration system 100.

[0024] The front-end interface 120 enable us to present information to users 103 using standard presentation methodologies. The front-end interface also allows the user to modify information and to generate transactions to initiate enterprise processes or workflow. The front-end interface can be modified to meet changing requirements of the enterprise.

[0025] The middle tier 130 uses our mobile agents 101 to provide an infrastructure for highly flexible, robust and scaleable distributed applications. The middle tier combines server technology with a customizable business rules engine and an application framework. The middle tier also provides for the deployment of disconnected applications for mobile users. That is, the middle tier allows the mobile user to perform tasks while disconnected from the system 100.

[0026] The design tools 140 support the definition of XML document formats. The design tools also allow us to define mappings of the XML document formats and the legacy data formats, and to provide for the automated generation of forms for user presentation via the front-end interface. These components are now described in greater detail.

### Back-End Interface

[0027] The back-end interface 110 is composed of one or more service bridges 112. The service bridges provide highly efficient access to various legacy systems. Hereinafter, we will frequently call the legacy systems "data sources" 111. We do not care how the legacy systems are programmed, or how their applications are structured. That is, the back-end interface of our integration system provides a generic and uniform access interface to the highly diverse legacy systems without requiring special knowledge of internal, legacy interfaces of the linked systems.

[0028] Semantically, we model the back-end interface

as an XML document publishing and management system. We see the data source as "publishing or "serving" XML documents containing enterprise information. The back-end allows users to add, update, delete, browse, and search for documents in the data source. We chose this semantic model of interaction because it provides a generic interface through which many disparate legacy systems can be accessed.

[0029] A particular data source 111 can manage multiple types of documents, such as customer accounts, purchase orders, work items, work lists, and the like. Any document in any data source can be uniquely identified and retrieved by a document identification (id) 104. In our implementation, and keeping within the spirit of XML, we use a document identification 104 that is conceptually similar to a Web page Universal Resource Locator (URL), although different in detail. As shown, the service bridges include a bridge framework (BF) 113 and a data source-specific runtime access component (RAC) 114. The service bridge is described in greater detail below with reference to Figures 4-9.

### Bridge Framework

[0030] The bridge framework 113 provides generic high level access services for the back-end interface. The framework is relatively independent from the specifics of the linked legacy systems and is implemented with reusable code. The bridge framework performs user authentication, and identifies the user making a request of the data source. The bridge framework also identifies agents 101 making requests, and provides a means to map a generic user identity to specific "logon" information required by any of the legacy data sources, e.g., a username and a password. The bridge framework operates securely such that any sensitive data-source logon information, such as a clear-text password, is encrypted.

### Connection Pooling and Document Management

[0031] The framework also manages objects involved in establishing and maintaining a connection to the data source, and provides for connection sharing and pooling. Connection pooling and sharing is used when the establishment of a connection or session with the data source is too expensive to perform on a per user basis. The connection pooling and sharing mechanism is based on "user groups." All members of a user group access a particular data source via a shared connection pool. The connections in this pool are established within the user context of a "pseudo-user account."

[0032] A pseudo-user account is a special data source account that represents a group of users instead of an individual user. Thus, if we have two user names, "john@accounting" and "jim@accounting," the two accounting users both access the data source within the context of the accounting pseudo user account. Con-

nection pooling may not be necessary for all back-end data sources, but certainly is required for relational database access.

### Document Caching

**[0033]** The bridge framework also provides a tunable caching facility to increase system performance. As stated above, a primary function of the back-end interface is to access legacy data and convert that data into the XML format. The bridge framework maintains XML documents in a cache 115 so that a subsequent request to retrieve the same data can bypass any data access or conversion work overhead by accessing the cached XML document.

**[0034]** The caching in our system is tunable. For a given type of document, a system administrator can specify caching parameters 116 such as whether caching should be enabled, a maximum lifetime before cache entries become stale, a maximum cache size, whether the cache 115 should be a persisted disk and re-used at next server startup. For document types that contain highly volatile data, caching can be disabled or cache entries can be set to expire quickly. For documents containing data that changes rarely, the caching parameters can be set aggressively to retain the documents in the cache.

### Runtime Access Component

**[0035]** The runtime access component (RAC) 114 is specific for a particular data source 111. The RAC uses application programming interfaces (APIs) and structures of the legacy data source to access the data and to map the data into the XML format. The exact semantics of how the data are mapped to the XML format vary. For example, the mapping can be for widely used legacy databases, such as, JDBC, JDBT, SAP, or SQL. An example JDBC implementation is described below with reference to Figure 4. The RAC supports the following database access operations.

#### Query

**[0036]** The "query" operation retrieves a document from the data source. The caller supplies the *id* 104 of the document to fetch. The bridge service returns the specified information in the form of a XML document according to one of the standard programming models supported by W3C, for example, a DOM document object or a SAX document. DOM (Document Object Model), is a programming interface specification that specifies a tree which applications may then explore or modify. SAX is an event-based tool, more or less 'reading' the document to the application using a set of named methods to indicate document parts. SAX is typically used where efficiency and low overhead are paramount, while the DOM is used in cases where applications need

random access to a stable tree of elements. The interface allows us to generate and modify XML documents as full-fledged objects. Such documents are able to have their contents and data "hidden" within the object, helping us to ensure control over who can manipulate the document. Document objects can carry object-oriented procedures called methods.

**[0037]** In the case of a relational database, the query operation maps to a SQL SELECT statement with a WHERE clause specifying which record or records from the database are contain in the document.

#### Update

**[0038]** The "update" operation modifies existing data in the legacy data source. The caller supplies the *id* of the document and a XML document containing only the fields to be modified. In the case of the relational database, the update operation maps to a SQL UPDATE statement.

#### Delete

**[0039]** The "delete" operation removes a document from the data source. The caller supplies the *id* of the document to delete. In the case of the relational database, the delete operation maps to a SQL DELETE statement.

#### Add

**[0040]** The "add" operation inserts a new document into the data source. The caller supplies the document in the form of a DOM Document object. The bridge service returns the *id* of the newly added document. In the case of a relational database, the add operation maps to a SQL INSERT INTO statement.

#### Browse

**[0041]** The browse operation, also known as "buffering," browses all of the documents in the data source of a certain type. The caller supplies the type of document to browse. The bridge service returns a browse object similar to a JDBC result set. The browse object allows the caller to traverse the results in either direction, jumping to the first or last document, and to re-initiate the browse operation. In the case of a relational database, the browse operation maps to a SQL SELECT statement that returns multiple records.

#### Search

**[0042]** The search operation browses the data source for all documents of a certain type that meet a pre-defined search criteria. The search criteria can be a list of fields and values which the caller wants to match against records in the database. For example, the caller might

request all customer records that contain a "state" field matching the string "MA." The caller supplies the type of document to browse as well as a document containing the fields to be matched. The bridge service returns a browse object as above. In the case of a relational database, the search operation maps to a SQL SELECT statement in which the WHERE clause contains the LIKE operator.

#### Front-End Interface

[0043] The front-end interface 120 is responsible for user presentation and interaction. The front-end interface uses "forms" to allow users to view and modify information. As an advantage, the front-end interface provides a "thin" user interface, with simple interactivity that can easily be customized as the environment in the enterprise changes. The front-end forms use HTML 121, HTTP 122, Javascript, Java servlets 123, Java applets, and plug-ins as necessary. Being Web based, the user 103 can use any standard browser 124 to interact with the system from anywhere there is an Internet access point.

#### HTTP Communications

[0044] The HTTP is used as the communication mechanism between agents and users. The user 103 browses and modifies information, and initiates processes via the web browser 124. User requests are routed to agents 101 via HTTP and through the Java servlet. The servlet 123 in turn communicates with a front-end service bridge 125 that serves as an interface for the agents 101.

[0045] The servlet/service bridge combination 123/124 supports the establishment of user sessions that are the channel for two-way communication between the user and the agents. Within the context of a session, the user can send HTTP GET or POST requests to the agents, and the agents process such requests, and send back an HTTP response. Sessions allow the user to wait for an agent to arrive and allow an agent to wait for a user to connect.

#### HTML Form Style Sheets

[0046] We accomplish the display of information to users with HTML, web pages, and web forms. As stated above, the information that agents retrieve from data sources is in the form of the XML documents 102. To format the XML documents into a form suitable for users, the front-end servlet 123 converts the XML document to a HTML page using a style sheet 126, e.g. XSL, JSP or some other data replacement technique as described below. The result of this conversion is the HTML page containing the information in a user-friendly format. By applying the style sheet, the servlet recognizes and replaces certain data from the XML document and

converts the data to HTML form.

[0047] For example, a particular XML document 102 includes the following information:

```
<customer>
    <firstname>John</firstname>
    <lastname>Smith</lastname>
</customer>
```

[0048] The HTML style sheet 126 for this document is as follows:

```
<html>
    <h1>'customer.firstname'</h1>
    <h2>'customer.lastname'</h2>
</html>
```

[0049] After applying the style sheet to the XML document, the resultant HTML form 121 would appear as:

```
<html>
    <h1>John</h1>
    <h2>Smith</h2>
</html>
```

[0050] The style sheet supports accessing all of the elements and attributes in the XML documents, and iteration over groups of repeating elements.

[0051] For example, an XML document contains:

```
<customer type="preferred">
    <firstname>John</firstname>
    <lastname>Smith</lastname>
</customer>
```

[0052] The "type" attribute of the customer is accessed by using a syntax such as the following:

```
'customer.attr[type]'
```

which yields the value "preferred." Given a document containing repeating groups as follows:

```
<customers>
    <customer type="preferred">
        <lastname>Smith</lastname> </customer>
    <customer type="standard">
        <lastname>Jones</lastname>
    </customer>
```

[0053] The "lastname" element of the second customer is accessed using a syntax such a 'customer[1].lastname' which yields the value "Jones." To iterate over all of the customers and access their "type" attributes, an expression such as:

```
'iterate(i=customers.customer) {
    i.attr[type]
```

can be used to produce first the string "preferred," and then "standard."

#### Validation

[0054] The front-end interface also supports the validation of user entered information. Field validation information supplies some immediate feedback and interactivity to the user. Field validation also increases application efficiency by detecting common errors within the web browser process before any other network traffic is incurred or application logic is executed. Client side validation can be broken down into two related levels.

[0055] The front-end interface also supports the validation of user entered information. Field validation information supplies some immediate feedback and interactivity to the user. Field validation also increases application efficiency by detecting common errors within the web browser process before any other network traffic is incurred or application logic is executed. Client side validation can be broken down into two related levels.

## Field-Level

[0055] Field-level validation performs simple checks on user entered data to validate that the information is of the correct format or data type. For example, field-level validation can validate that a user enters numeric values in a particular field, or uses a proper date format. We implement field-level validations with Javascript. A library of common validations is supplied as a script file on a web server. The library has a ".js" file extension. This script file can be included into HTML forms as desired using the <script> HTML tag. Validation is enabled for a field by indicating the name of an appropriate validation routine, e.g. "onChange," within an event handler of the field. The event handler is triggered when an INPUT field changes. Setting up validation for a field requires HTML coding as follows:

```
<input type="text" name="birthdate" onChange="validateDate(birthdate)">
```

The validation library provides routines for common data types such as dates, times, currency, etc. The validation library can also provide a pattern matching ability allowing user input to be matched against arbitrary patterns, e.g., a pattern \$##.## to match a monetary amount.

## Cross-Field Validation

[0056] Cross-field validation allows for more complex validations. In this type of validation, the contents of one field depends on the contents of another field. For example, cross-field validation can detect a situation where a telephone number must be entered. Such validation usually requires a more detailed knowledge of the requirements of the application.

## Middle Tier

[0057] The middle tier 130 provides the "glue" that links the back-end and the front-end interfaces. The middle tier utilizes the mobile agents 101 to communicate with the interfaces. The middle tier also provides support for disconnected applications and users. In addition, the middle tier customizes the system 100 to the needs of specific enterprise functions without actually having to reprogram the legacy systems.

[0058] The middle tier supports the automation of complex workflow and complex validations of data that may require access to multiple data sources. As a feature, the middle tier uses a rules engine (RE) 131 operating on rules stored in a database 132. The rules are defined in a rules language, and can be retrieved by the agents 101 as needed.

[0059] In a typical scenario, the user launches an agent 105 due to interaction with the browser 124. The agent carries an XML document, e.g., a purchase order 106, to the rules database 132. The agent retrieves the appropriate rule for processing the order, such as a purchase order workflow. The agent then interprets the rule

to appropriately route the document to the locations in the network specified by the rule. The rule can include a travel itinerary, as well as instructions on how to interact with the data sources.

5 [0060] As an advantage, the operation of our system is always current. As rules change so does the operation of the system. The agents always execute according the current state of the rules database.

## 10 Design Tools

[0061] As shown in Figure 2, the primary purpose of the design tools 140 is to generate 141 XML document type definitions (DTD) 142, to specify 143 data mappings, i.e., RACs 114, to encode 144 rules 132, and to design 145 user interfaces 126.

### Document Type Definitions

20 [0062] The step 141 identifies the different types of document information (DTD) 142 that needs to be shared by the various data sources 111 of the back-end 110 and the browser 124 of the front-end 120. This information is specified in the DTDs. For example, to share purchase order information between systems, the type of information needed in a purchase order needs to be identified, then that information needs to be encoded in a corresponding DTD. In one embodiment, the design tools use the service bridge to extract schemas 25 from the data sources.

### Data Mapping

30 [0063] After a data source independent data format 35 has been generated, the mappings between the XML format and legacy formats for a particular database needs to be specified as shown in Figure 3. A query operation to a relational databases 111 involves extracting the schema of the database by generating a SQL runtime access component (RAC) 114 which makes the JDBC calls to the database, converting the resulting data 40 into the XML format, and handing the XML document 113 to an agent 101. The access components can be implemented as Java code. The agent delivers the XML 45 to the front-end 120 for conversion to the HTML form 121 using the style sheet 126 so that the data can be viewed by the user 103 using a standard browser 124.

50 [0064] Conversely, the update operation converts the HTML form to the corresponding XML document. The XML document is converted to a legacy format and the RAC modifies the data source using its schema. For other legacy data sources that are not specified by a schema or some other metadata, the mapping may need to be done by means that access the APIs directly.

### 55 Rule Encoding

[0065] After the data format definition is generated,

and the RAC has been specified to access the appropriate data source, the next step is to encode what agents are going to do with the information. In a simple data replication system, an agent may retrieve modified records from a master database, travel to the location of a backup database, and then update the backup database with a copy of the modified record. This process involves the encoding of a specific rule.

### Designing the User Interface

**[0066]** As shown in Figure 2, generating the user interface requires three steps: manipulating document type definitions (DTD) 145, importing DTD 146, and generating DTD from database schema 147.

### Authoring DTD

**[0067]** The design tools 140 allow the system designer to define, design, and manipulate XML and HTML DTDs. A DTD 142 defines the name of the following document elements: the contents model of each element, how often and in which order elements can appear, if start or end tags can be omitted, the possible presence of attributes and their default values, and the names of the entities.

**[0068]** Because the DTDs represent many different types of documents in the system, this step essentially defines the data types of the enterprise's computerized applications. As an advantage, the resulting DTDs do not directly tie the system to any specific legacy data source, nor do the definitions preclude the integration of other legacy systems in the future.

### DTD Import

**[0069]** The tools also allow one to import already existing DTD definitions. Such functionality can be used in environments where DTDs have already been defined for standard document types. These DTDs may have been defined by standards bodies or a designer of the legacy system.

### DTD generation from Database Schema

**[0070]** This part of the tools automatically generate DTDs from existing database schema.

### XML ↔ SQL Mapping Definition

**[0071]** Given the existence of the DTDs, the system 100 provides tools that map between legacy back-end data formats and XML document formats. In the case of relational database access, these mappings link tables, columns, and fields from the legacy database to elements and attributes of the XML documents as defined by the DTDs. This also allows the definition of several distinct mappings, each of which involves accessing

slightly different information in the data source.

### Data Mappings

#### 5 Query Mapping

**[0072]** A query mapping enables an agent to retrieve information from a legacy data source. In the case of a relational database, this mapping specifies the contents 10 of the SELECT statement, including any information relevant for a table join. A query mapping for a purchase order may involve accessing a purchase order table, a customer table, and a product catalog table.

#### 15 Update Mapping

**[0073]** An update mapping allows an agent to modify information in the data source. This involves specifying the contents of an UPDATE statement. An update mapping for a purchase order involves updating the purchase order table, but not modifying the customer table or the product catalog table.

#### Delete Mapping

**[0074]** A delete mapping allows an agent to delete information in the data source. This involves specifying the contents of a DELETE statement. A delete mapping for a purchase order involves deleting a record or 30 records from the purchase order table, but not modifying the customer table or the product catalog table.

#### Add/Create Mapping

**[0075]** An add/create mapping allows an agent to add information to the data source. This involves specifying the contents of an INSERT statement. An insert mapping for a purchase order involves adding a record or records to the purchase order table, but not modifying 40 the customer table or the product catalog table.

### Schema Extraction and Caching

**[0076]** In order to allow for mapping between a legacy 45 database schema and XML DTD formats, the mapping design tool extracts the schema from legacy databases. Because schema extraction is an expensive and time consuming task, the tools allow one to save extracted schemas on a disk for subsequent use.

#### 50 Form Generation

**[0077]** The tools will also allow one to automatically 55 generate a form from a DTD. Such a form may require minor modifications to enhance the physical appearance of the form. For example, color or font size of text can be adjusted to enhance usability.

### Embedding Binary Data In XML Documents

[0078] Some enterprise applications may need to retrieve arbitrary binary data from the data source 111. For example, a legacy database contains employee information. Included with that information is a picture of the employee in standard JPEG format. The employee information is stored as a single table named "employees," which has a schema as Table 1, where the field <image> represents the picture:

Table 1

ID	Name	HireDate	Photo
1	John Smith	1/1/96	<image>

[0079] The XML document that retrieves the above table appears as follows:

```
<employee>
  <ID>1</ID>
  <name>john Smith</name>
  <hiredate>1996-29</hiredate>
</employee>
```

XML, by itself, does not naturally lend itself to the inclusion of binary data. To deliver this information for display in a web page, the service bridge 112 could encode the SQL record in an XML document as follows:

```
<employee>
  <ID>1</ID>
  <name>john Smith</name>
  <hiredate>1996-29</hiredate>
  <Photo href="http://server/directory/john.
jpeg"/>
</employee>
```

[0080] However, there are a number of problems with this type of approach. First, it is the responsibility of the user to issue the proper additional commands to retrieve the linked document before it can be displayed, e.g., the user must click on the URL of the picture. Second, the DTD for the XML document must specify the URL. For most legacy databases, it is unlikely that the records storing the binary data are accessible via an HTTP URL. Furthermore, the binary data is transported through the system by a follow on transport, such as HTTP. For reliability, security, consistency, and other reasons we prefer to carry all data, including binary data with the agents.

[0081] To allow the servlet 123 to generate an agent that can access the binary data, we define a new type of URL. The new URL incorporates the location of the binary data, as well as a unique "name" that can be used to retrieve the binary data. The URL contains the host-name of the data source, a service name, an action name that can be used to perform the retrieval of the binary data, and a document identification referring to the binary data. This still results in a fairly complex URL.

[0082] Using multiple requests to retrieve the binary data is inconsistent with our agent model. Agents try to

use the network effectively by batching data into fairly large self-contained packets. This is very different than the hypertext model used on the web in which a single page display can lead to multiple network requests.

5

### Compound Documents

[0083] In an alternative solution, we define a compound document. In a compound document, the binary data is embedded in the same document as the textual XML data. This approach is consistent with our agent driven system that attempts to transport data as larger batches. Compound documents can be built in two ways.

15

### Embed Binary Data into XML Text Element

[0084] The binary data is embedded directly into an XML text element. This can be done as long as the binary data is encoded in such a way that the data only contain XML characters. Such an encoding could be based on the Base64 encoding. With Base64, special characters, such as "<" and ">," are replaced with equivalent entities (i.e., &lt; and &gt;). We also can use a character data (CDATA) section to work around the problem of illegal characters within the Base64-encoded data. We may want to prefix the embedded binary data with standard mime headers that specify content type, encoding, and name. Such a format for the photo element appears as follows:

```
<Photo>
  Content-Type: image/jpeg
  Content-Encoding: base64
  Content-Name: john.jpeg
  9j/4AAQSKZJ.....gEASABIAAD/
</Photo>
```

[0085] It should be noted that this alternative increases the size of the binary data by 33% as well as increasing the overhead to encode and decode the data.

[0086] This alternative requires that a SQL RAC extracts the binary data and encodes the data into Base64, and then adds the encoded data to the XML document with the proper mime headers.

### Compound Document Encoded as Mlme Document

[0087] Another alternative, embeds both the XML document and the binary data into separate parts of a multipart mime document. Each part of the overall document has a Content-ID which is referenced from a standard XML link, in part, such a format appears as follows:

```
Content-Type: multipart/related; boundary= "--XXXXX"
--XXXXX
Content-Type: text/xml
Content-ID: doc
  <Photo href="cid:photo"/>
--XXXXX
```

Content-Type: image/jpeg  
 Content-Encoding: base64  
 Content-Name: john.jpeg  
 Content-ID: photo  
 9j/4AAQSKZJ... gEASABIAAD/  
 ---XXXX---

**[0088]** With this alternative, the binary data may not need to be encoded. However, this requires that agents also retrieve MIME documents via the RAC.

### JDBC Service Bridge

**[0089]** Figure 4 shows details of a preferred embodiment of a service bridge 400 of the back-end interface 110 for accessing a data source. In this embodiment, JDBC is used to access a SQL type of database. The bridge 400 includes a public interface 410, JDBC runtime access component (RAC) 420, XML-SQL data mapping 430, and a document cache 440 as its main components.

### Public Interface

**[0090]** As stated above, the public interface 410 provides the means by which agents access the data sources 111. The public interface allows data retrieval, modification, and addition. As an advantage, the public interface 410 makes no assumptions about how data in the legacy database 111 is sourced or maintained. Instead, we make the public interface resemble the GET/PUT model of HTTP.

### JDBC Run-Time Access Component

**[0091]** The JDBC access component 420 is responsible for establishing and managing JDBC connections, building and executing SQL statements, and traversing result sets. This component works entirely within the context of JDBC and SQL.

### XML-SQL Data Mapping

**[0092]** The XML-SQL data mapping 430 uses the mapping information generated by the design tools 140 to map data between XML and SQL.

### Document Cache

**[0093]** The document cache 440 operates entirely with XML documents. XML documents that have been retrieved from the data source can be cached for fast future retrieval. The caching services are configurable so that maximum cache sizes and cache item expiration times can be specified. Caching can be disabled for certain classes of documents which contain highly volatile information.

**[0094]** Figure 5 shows the public interface 410 in greater detail. The interface supports four basic types

of accesses, namely get 510, put 520, add 530, and delete 540.

**[0095]** At the heart of the interface is the document *id* 104. The document *id* is a string which uniquely identifies every document instance within the data source. The document *id* can be thought of as corresponding to the URL of a World Wide Web document, or to the primary key of a record in a database. Although the *id* has a different format than a URL, it does serve as a document locator.

**[0096]** In order to interact with information in the legacy data source, an agent needs to provide the *id* for the document containing the information. The *id* contains multiple sections of information and follows the following pattern.

**[0097]** The first character of the *id* string specifies a separator character (S) 501 that is used to separate the different sections that make up the document *id*, e.g., a colon (:). This character is used in conjunction with a Java StringTokenizer to parse the document *id*. The subsequent information in the *id* includes name=value pairs (N, V) 502. One pairs 502 specifies a document type, e.g., ":type=custom\_list:"

**[0098]** In most common cases, the *id* 104 also contains a key specifying the exact document instance in order to uniquely identify an individual document in a data source. For example, in a document containing customer information, this key contains a data source specific customer number or a customer id. Within the service bridge, this key is mapped to a WHERE clause of a SQL statement. For example, an agent can request customer information for a particular customer by specifying an *id* string as follows:

"":type=customer:key=SMITH:".

This request results in a SQL query to the database that appears as follows:

SELECT \* FROM Customers WHERE Customers.ID=SMITH

The exact semantics of how they key is mapped into the resultant SQL statement is specified by the design tools 140.

**[0099]** The key portion of the *id* can be composed of multiple pieces of information separated by, for example, commas. Such a key is used in cases in which the WHERE clause of the corresponding SQL query needs multiple pieces of information to be specified by the agent. An example of this is a document containing a list of customers, where the customers names are within a certain alphabetic range, for example, "all customers whose last names begin with the letters A or B. Such a document has an *id* as follows:

"":type=custom\_list\_by\_name:key=A,Bzzzz:"

In this case, the request would map into a SQL statement resembling the following:

SELECT \* FROM Customers  
 WHERE Customers.LastName BETWEEN A, Bzzz

10

## Implementation Details of the Service Bridge

### Database Access

### User Authentication

**[0100]** The service bridge is responsible for performing any authentication necessary in order to establish a database connection. This may involve supplying a database specific username and password or other login information. When a database access (get, put, add, delete) is made by an agent, the bridge examines the agent's runtime context to determine the user identity associated with the agent.

**[0101]** After the agent's identity has been ascertained, the service bridge maps the identity into simultaneous database-specific user identification using a mapping table generated by the design tools. For example, the mapping maps the user identity "steve@accounting" into an Oracle-username "steve."

**[0102]** In order to establish a connection to a database on behalf of a user, the service bridge retrieves both the username and clear-text password for the corresponding database user account. In such cases, the clear-text password is stored in the identity-mapping table. For security reasons, the table is encrypted on disk using a public/private key pair.

### Connection Management

**[0103]** To enhance performance and scalability, the service bridge supports database connection pools. This means that multiple users share a common pool of JDBC connections. Establishing a database connection can be a slow and relatively expensive operation. The use of shared connection pools decreases this expense.

**[0104]** The basis for this connection sharing are "users groups." When an agent attempts an operation which requires a connection to a database, the service bridge performs that operation using a connection established in the context of a special "pseudo-user" account. The pseudo-user is a database system account that represents not an individual user, but instead a particular group of users. A pool of such pseudo-user connections is available for use by all of the agents of the group. The service bridge generates and maintains a connection pool for each distinct group of users who access the bridge.

**[0105]** Figure 6 shows agents 101 for three users *tom*, *joe* and *david* 601-603 accessing the data source 111. Two of the users, *tom@users* and *joe@users*, are members of a users group. The third user, *david@managers*, is a member of a "managers" group. When these agents attempt to access the database, the two members of the users group share a connection pool 610 that was established with the credentials of the "users" pseudo-user. The third agent will communicate with the database using a separate connection pool 620 established with

the credentials of the "managers" pseudo-user.

**[0106]** A connection pool for a particular group is generated when a member of the group makes the first access request. Connections within the pool are constructed as needed. The service bridge does not pre-allocate connections. After a configurable, and perhaps long period of inactivity, the connection pool is closed to free database resources. If a connection pool for a particular group has been closed due to inactivity, then any subsequent request by a member of that group results in the generation of a new pool. When a request is completed, the connection allocated for that request is returned to the pool. A maximum number of connections in a pool can be specified. If no connections are available when a request is made, then the request is blocked until a connection becomes available.

### Statement Construction and Execution

**[0107]** The actual generation and execution of SQL statements is performed by a separate "modeler" object. The modeler object is generated by the design tools 140. For each type of document used in the system, there is a distinct modeler object. Each modeler knows how to construct exactly one type of document. During the design process, one specifies what information is to be retrieved from the database, and how to map the information into an XML document. The design tools serialize and save the modeler objects in a ".ser" file. At runtime, the service bridge loads and de-serializes the modeler objects from the ".ser" file. The resultant modeler objects are able to perform all of the data access and mapping functions required to retrieve information from the data sources. As stated above, SQL to XML data mapping is performed by the modeler object designed for a particular document type.

### Data Caching

**[0108]** To improve the performance of document retrieval, the data service caches database information as converted XML documents. When a first request is made to retrieve a document, the service performs the SQL access and SQL to XML data mapping as described above. The resultant XML document is added to the cache of documents 440 maintained by the service bridge. Any subsequent request to retrieve the document will be satisfied by retrieving the document from the cache, bypassing the need for an additional expensive database access and mapping.

**[0109]** When an update or addition is made to a data source, the cache is updated to reflect the new information. The update to the cache is made only after the SQL statement performing the update of the end database has been completed successfully. This prevents the cache from storing information that has not been committed to the database due to errors or to security restrictions.

[0110] The XML document cache is configurable to specify a maximum size of the cache, the maximum amount of time a single document can be retained in the cache before it becomes stale, and whether the cache should be persisted to disk, in which case the cache can be re-used after a server restart. One can also customize how different classes of documents are cached. If a document represents highly volatile information, then caching can be disabled for that class of document. If a document class is completely (or virtually) static, then documents of that class can be cached for a very long time.

### Execution Flow

[0111] The following section describes the execution flow for basic database access requests. Figure 7 shows the steps 700 of a "get" or retrieval access in greater detail. After the request is received from the agent 710, the caller and document identity are determined 720, 730. The group specific cache is identified 740, and the cache is checked 750. If the cache stores the document, return the document in step 755. Otherwise, locate the XML-SQL mapping 760, construct the select SQL select statement 770, retrieve the connection 775, and execute the statement in step 780. Next, the result set is "walked" 785, fields are extracted 790 to build the XML document 794, the document is cached 796 and returned to the agent in step 798. Figure 8 shows the steps 800 for the addition (add) and modification (put) similar to the get steps. The delete request simply deletes data from the database as shown at 540 in Figure 5.

### Run-time Object Hierarchy

[0112] Figure 9 shows the run-time hierarchy 900 of objects of the service bridge 110. The objects can be classified as data source independent 901, and data source dependent 902. The data source independent object 901 includes data source factory object 910 indexed by group name, group specific data source objects 920, document factory objects 930 (one per document), document cache objects 940, document builder objects 950, connection pool objects 960, mapping table objects 970, document manager objects 980, and the data source manager objects 990. The data source dependent object 902 include source connection 991, string authentication 992, document map 993, and specific driver objects 994.

[0113] Although the invention has been described by way of examples of preferred embodiments, it is to be understood that various other adaptations and modifications may be made within the spirit and scope of the invention. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

### Claims

1. An enterprise integration system coupled to a plurality of data sources, the plurality of data sources using different data formats and different access methods, comprising:
  - 5 a back-end interface configured to convert input data source information to input XML documents and to convert output XML documents to output data source information;
  - 10 a front-end interface including means for converting the output XML documents to output HTML forms and for converting input HTML forms to the XML documents;
  - 15 a middle tier including a rules engine and a rules database;
  - 20 design tools for defining the conversion and the XML documents;
  - 25 a network coupling the back-end interface, the front-end interface, the middle tier, the design tools, and the data sources;
  - 30 a plurality of agents configured to communicate the XML documents over the network and to process the XML documents according to the rules.
2. The system of claim 1 wherein the back-end interface further comprises:
  - 35 a public interface;
  - 40 a document cache; and
  - 45 a run-time access component.
3. The system of claim 2 wherein the public interface forwards the input XML document to the plurality of the agents for distribution, and the public receives the output XML documents for storing in the plurality of data sources.
4. The system of claim 2 wherein the document cache includes caching parameters.
5. The system of claim 2 wherein the caching parameters include a maximum lifetime for each cache entries, a maximum cache size, and a persistency indicator.
6. The system of claim 2 wherein the run-time access component generates access requests for the plurality of data sources.
7. The system of claim 6 wherein the access requests include query, update, delete, add, browse, and search.
- 55 8. The system of claim 1 wherein the XML documents include binary data.

9. The system of claim 8 wherein the binary data is referenced by a Universal Resource Locator.

10. The system of claim 8 wherein the binary data is embedded as a compound document. 5

11. The system of claim 10 wherein the compound document embeds the binary data as an encoding in a character set. 10

12. The system of claim 10 wherein the compound document embeds the binary as a MIME document.

13. The system of claim 1 wherein the input documents are presented to a browser. 15

14. The system of claim 1 wherein each XML document is identified by a document identification.

15. The system of claim 14 wherein the document identification is a character string. 20

16. The system of claim 15 wherein the character string includes a plurality of sections, and a first character of the string is a section separator. 25

17. The system of claim 16 wherein one of the sections stores a document type.

18. The system of claim 15 wherein one of the sections stores a key to an instance of the XML document in one of the data sources. 30

19. The system of claim 1 wherein the back-end interface performs user authentication. 35

20. The system of claim 1 wherein the back-end interface supports database connection pools.

21. A method for integrating a plurality of data sources, the plurality of data sources using different data formats and different access methods, comprising:  
converting input data source information to input XML documents and converting output XML documents to output data source information; 45  
converting the input XML documents to input HTML forms and converting output HTML forms to the output XML documents; 50  
providing a rules engine and a rules database; defining the converting and the XML documents;  
communicating the XML documents over a network using agents; and 55  
processing the XML documents by the agents according to the rules database.




FIG. 1a  
PRIOR ART




FIG. 1b




FIG. 2




FIG. 3





FIG. 4






FIG. 6




FIG. 7




FIG. 8




FIG. 9



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 01 10 2032

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int.Cl.)						
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim							
P,X	EP 1 030 254 A (SD & M) 23 August 2000 (2000-08-23) * abstract; figure 1 *	21	G06F17/30						
A	"Microsoft Announces Finalized BizTalk Framework" MICROSOFT PRESSPASS, [Online] 6 December 1999 (1999-12-06), XP002166805 Redmont, Wash. Retrieved from the Internet: <URL: http://www.microsoft.com/PressPass/pr ess/1999/Dec99/BIZTALK1.0pr.asp> [retrieved on 2001-05-08] * the whole document *	1,21							
A	VAN VOOREN L: "XML and legacy data conversion: introducing "consumable documents"" SGML EUROPE '97. CONFERENCE PROCEEDINGS, PROCEEDINGS OF SGML '97. THE NEXT DECADE - PUSHING THE ENVELOPE, BARCELONA, SPAIN, 13-15 MAY 1997, [Online] pages 185-187, XP002166806 1996, Alexandria, VA, USA, Graphic Commun. Assoc, USA Retrieved from the Internet: <URL: http://www.infoloom.com/gcaconfs/WEB/ barcelona97/vooren57.HTM> [retrieved on 2001-05-08] * the whole document *	1,21	TECHNICAL FIELDS SEARCHED (Int.Cl.) G06F						
A	WO 99 23584 A (FRIEDMAN MARK M ;IOTA IND LTD (IL); STERN YONATAN PESACH (IL)) 14 May 1999 (1999-05-14) * abstract * * page 2, line 21 - page 4, line 21 *	1,21							
<p>The present search report has been drawn up for all claims</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">Place of search</td> <td style="width: 33%;">Date of completion of the search</td> <td style="width: 34%;">Examiner</td> </tr> <tr> <td>BERLIN</td> <td>8 May 2001</td> <td>Triest, J</td> </tr> </table> <p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone  Y : particularly relevant if combined with another document of the same category  A : technological background  O : non-written disclosure  P : intermediate document</p> <p>T : theory or principle underlying the invention  E : earlier patent document, but published on, or after the filing date  D : document cited in the application  L : document cited for other reasons</p> <p>&amp; : member of the same patent family, corresponding document</p>				Place of search	Date of completion of the search	Examiner	BERLIN	8 May 2001	Triest, J
Place of search	Date of completion of the search	Examiner							
BERLIN	8 May 2001	Triest, J							

ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.

EP 01 10 2032

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

08-05-2001

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
EP 1030254	A	23-08-2000	NONE		
WO 9923584	A	14-05-1999	US	6161107 A	12-12-2000
			AU	1371599 A	24-05-1999